Spatial Vision, Vol. 5, No. 2, pp. 101–121 (1990)

Light–dark asymmetries in the Craik–Cornsweet–O'Brien illusion and a new model of brightness coding

BERNARD MOULDEN* and FRED KINGDOM†
Department of Psychology, University of Reading, Whiteknights, Reading RG6 2AL, UK

Received 18 July 1989; revised 1 June 1990; accepted 5 June 1990

Abstract—With the aid of a matching technique, the magnitude of induced brightness in bars bordered with Craik–Cornsweet–O'Brien (CCOB) edges was investigated as a function of the width and amplitude of those edges. Data were collected for stimuli with the sloping part of the edge on both the inside and outside of the bar, and also for stimuli with both positive-going and negative-going edges. The results confirmed previous reports that induced brightness was greater for CCOB stimuli with negative-going, as opposed to positive-going, edges and greater for CCOB stimuli whose edges contained outer, as opposed to inner gradients. A model of brightness coding is offered to provide an explanation for the specific anisotropies observed, as well as the general effects of stimulus amplitude and width on induced brightness. The model assumes that a symbolic description of brightness is generated separately from each of a number of different-sized 2DG (second difference of a Gaussian) filters, and the resulting brightness profile obtained by averaging across the separate descriptions. The ability of other brightness models to account for the data is also discussed.

INTRODUCTION

The familiar Craik–Cornsweet–O'Brien (CCOB) illusion has prompted a good deal of investigation over the last three decades (see review by Kingdom and Moulden, 1988), largely because it is a striking anomaly in the subjective appearance of surface reflectance that provides a challenge to any theory of brightness coding. The best known version of the illusion is one which was first demonstrated by Cornsweet (1970), and this is shown in Fig. 1.

The figure looks like a bipartite field, with the left half darker than the right, even though the two halves have identical reflectances except in the small region defined by the dividing contour. The critical feature of the contour, which we will refer to as the CCOB edge, is that it consists of a combination of both a sharp and gradual change in reflectance. Another version of the illusion is exemplified by the four classes of disc stimuli in Fig. 2. The CCOB edges in each disc consist of a sharp discontinuity bounded by a gradient on just one side (rather than having a gradient on both sides as in Fig. 1), like the first published version of the CCOB (O'Brien, 1958). Each of the four discs in Fig. 2 is characterized by a combination of two factors: (a) whether the CCOB edge is a negative- or a positive-going edge; and (b) whether the gradient is on the inside or the outside of the sharp discontinuity.

*Present address: Department of Psychology, University of Western Australia, Nedlands, Perth, WA 6009, Australia
†Present address: Department of Ophthalmology, Royal Victoria Hospital, 687 Pine Avenue West, Montreal, Quebec H3A 1A1, Canada

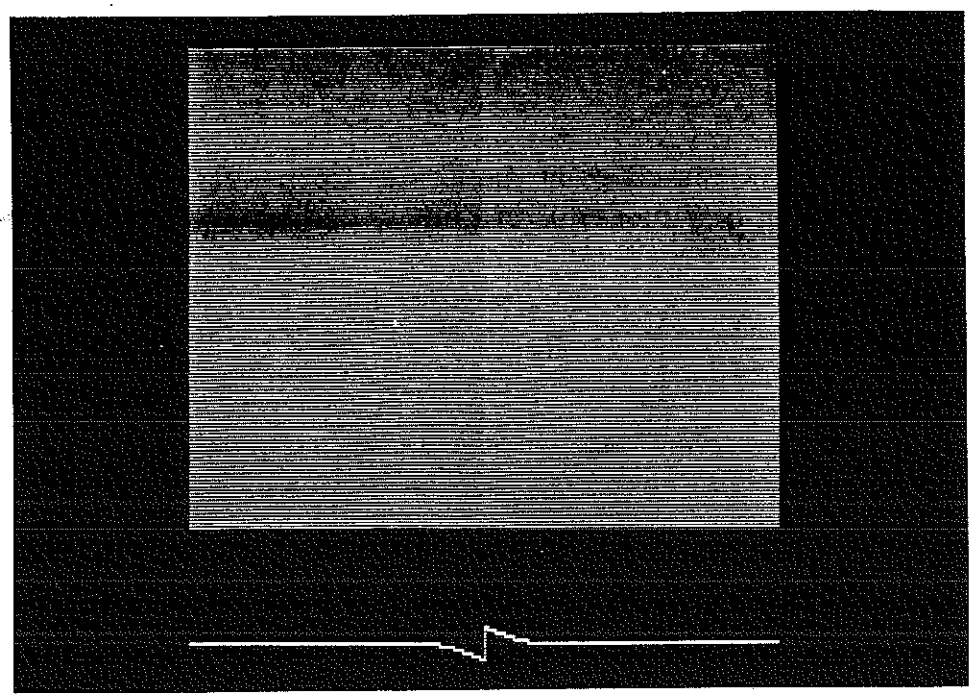

Figure 1. The Cornsweet illusion. The luminance profile of the stimulus is also shown.

The experiments described here were designed to measure the magnitude of induced brightness in the central region of one-dimensional versions of the stimuli shown in Fig. 2. The purpose of this communication is threefold: first, to show that the pattern of induced brightness is different for the four classes of stimuli shown in Fig. 2; second, to provide a computational model to account for those differences; and finally to compare that model with other current models of brightness coding.

There are previous reports of an anisotropy in the magnitude of induced brightness produced by positive and negative 'cusped' stimuli similar to those in Fig. 2. Hamada (1985) found that the magnitude of induced brightness was greater for negative than for positive cusped stimuli. He required his subjects to adjust the luminance of the regions internal to the cusps until they appeared equal in brightness to the outer regions. In this 'compensation method' the magnitude of induced brightness was given by the difference in luminance between the inside and outside areas at the PSE (point of subjective equality). He measured the magnitude of the effect as a function of the width of the CCOB edge for one amplitude (see Fig. 3) and found that while the magnitude of induced brightness increased at the same rate with gradient width for both positive and negative cusps, it did so by a greater constant amount for the latter. When gradient width was held constant and amplitude varied, induced brightness increased monotonically for the negative cusps, whereas for the positive cusps there was a small increase followed by a reversal—a reversal implying that the central region appeared darker than the surround.

Todorovic (1987) presented demonstrations similar to those in Fig. 2 and remarked that the magnitude of the effects in the stimuli with outer gradients appeared greater

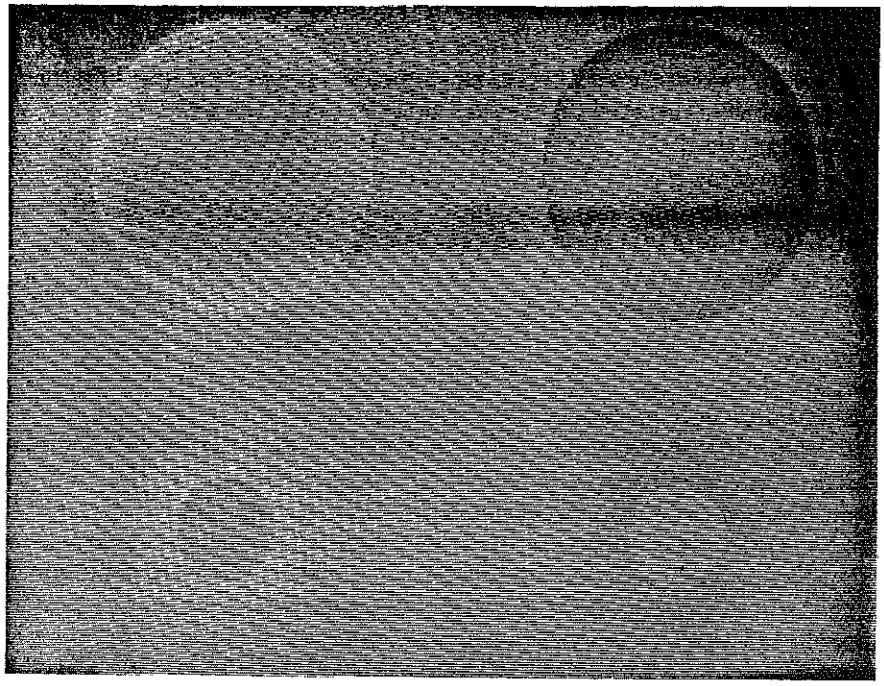

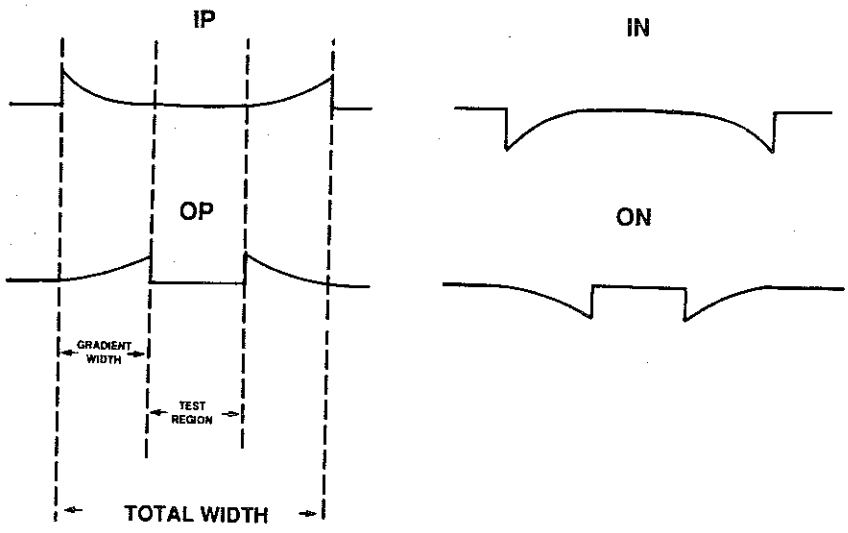

Figure 2. Four classes of CCOB discs. The stimuli are labelled as IP = inner positive-going; IN = inner negative-going; OP = outer positive-going; ON = outer negative-going. The top pair (IP and IN) have the sloping parts of the edge on the inside of the stimulus, while the bottom pair (OP and ON) have the sloping part of the edge on the outside of the stimulus. The left hand pair (IP and OP) have positive-going edges, while the right hand pair (IN and ON) have negative-going edges. The luminance profile of each stimulus is shown below, along with a description of the terminology employed to describe the spatial dimensions of the stimuli. For the actual experiments one-dimensional versions of these radial stimuli were employed.

than that in the stimuli with inner gradients. However, he reported no quantitative data.

We begin by presenting data on the magnitude of induced brightness for the four classes of pattern whose radial versions are shown in Fig. 2. We wished (a) to verify the findings of Hamada (1985) using a different technique and a greater range of gradient widths and amplitudes and (b) to provide quantitative data for CCOB patterns with outer gradients; the aim was to establish a data base which would provide the foundation for quantitative modelling.

METHOD

Stimulus generation

All stimuli were generated by an 8-bit Pluto II graphics display system interfaced to a Corvus Concept host computer. The programs were written in PASCAL using ASM68K assembly language subroutines to interface the host and graphics computers. The stimuli were displayed on a Barco type 2 TVMR monochrome TV monitor. The pixels on the screen were 0.69 mm in height by 0.33 mm in width subtending 2.03 by 0.98 arcmin respectively at the viewing distance of 114 cm. All luminance calibrations were performed with a purpose-built photodiode and amplifier system with a photometric filter. The aperture of the photodiode was positioned in front of the glass surface of the monitor and the photodiode was focussed onto a small patch of pixels, whose grey level had been preset.

Stimuli

We employed one-dimensional as opposed to radial stimuli. The CCOB edge luminance profiles were generated using a quarter cycle of a sinusoid:

$$L(x) = L_b + L_a + L_a \sin((2\pi/4W)x + p), \quad 0 < x < W, \tag{1}$$

where $L(x)$ is the luminance of the xth pixel, L_b is the background luminance of 20.0 cd/m^2, L_a is the luminance amplitude of the edge, W is the width of the gradient and p is the phase of the underlying waveform. For the left and right inner gradients, $p = \pi$ and $3\pi/2$ respectively, while for the left and right outer gradients, $p = 3\pi/2$ and π. Positive and negative stimuli were produced with respectively positive and negative values of L_a.

The stimuli all subtended 9.75 deg in height. The test region in the stimuli with CCOB edges was 1.6 deg in width. For the conditions with inner CCOB edges, three edge widths were employed, 0.65, 1.3 and 2.6 deg, while for the stimuli with outer CCOB edges only one width condition, 2.6 deg, was employed. The total stimulus width of the three width conditions was thus 2.25, 2.9 and 4.2 deg respectively. For all stimulus varieties six amplitude conditions were used, 0.5, 1.0, 2.0, 4.0, 8.0 and 16.0 cd/m^2 above (positive) and below (negative) the mean value of 20 cd/m^2. The match stimulus was a bar whose width matched the total width of the test stimulus and whose luminance was adjustable by the subject. The central test region of both the match and test stimuli was delineated by four short black vertical lines, each 20 arcmin in length, positioned in the four corners of the test region.

Subjects

The two authors acted as subjects. Both were experienced psychophysical observers, one (FK) with normal and the other (BM) with corrected vision.

Procedure

During an experimental session only one of the four classes of stimuli was displayed, but the various width and amplitude conditions were presented in random order. On each trial the test and match stimuli were alternated in time with a stimulus exposure duration and ISI (inter-stimulus-interval) of 1 s. During the ISI the screen was homogenous with a luminance of 20.0 cd/m^2. The task of the subject was to adjust the luminance of the match stimulus until its central region appeared equal in brightness to that of the test stimulus. For each condition each subject made a total of 10 measurements.

RESULTS

The results are displayed in Figs 3 and 4. In each figure, the results for the stimuli with both inner and outer gradients are presented together. The figures show the amplitude of the match bar (ordinate) required to match each stimulus as a function of the test stimulus amplitude (abscissa), for various gradient widths, and the two classes of stimulus (inner or outer gradient). Figure 3 presents the results for the positive stimuli, Fig. 4 those for the negative stimuli. It should be noted that the plots for the inner

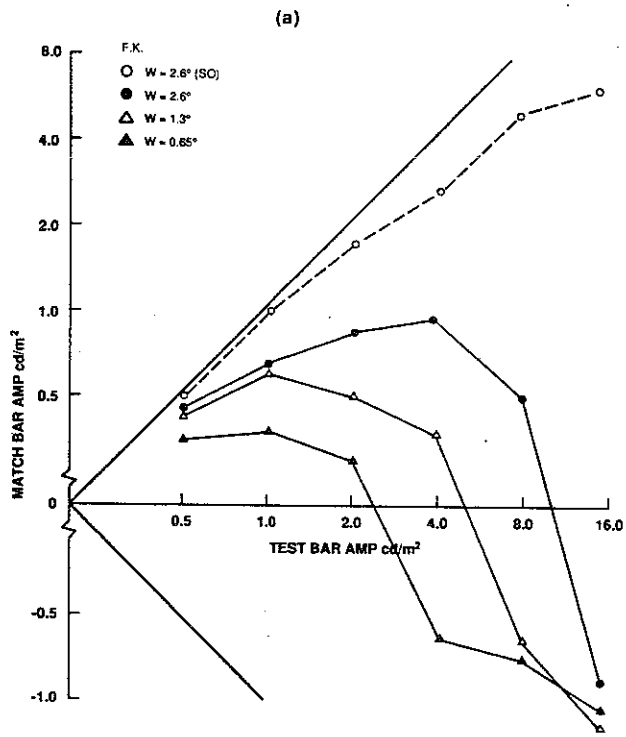

Figure 3. Results and model predictions for the stimuli with positive-going edges. The plots show the amplitude of the match bar required to make a brightness match with the central test area of the CCOB stimulus, as a function of the amplitude of the CCOB stimulus. Continuous lines = stimuli with inner gradients; dashed line = stimuli with outer gradients; W = width of edge. (a) subject FK, (b) subject BM. Positive values of test bar amplitude imply that the central test area appeared lightened, negative values that the central test area appeared darkened. In (c) the predictions from the model are shown for the conditions shown in (a) and (b). The units in (c) are arbitrary. (*Contd. on next page.*)

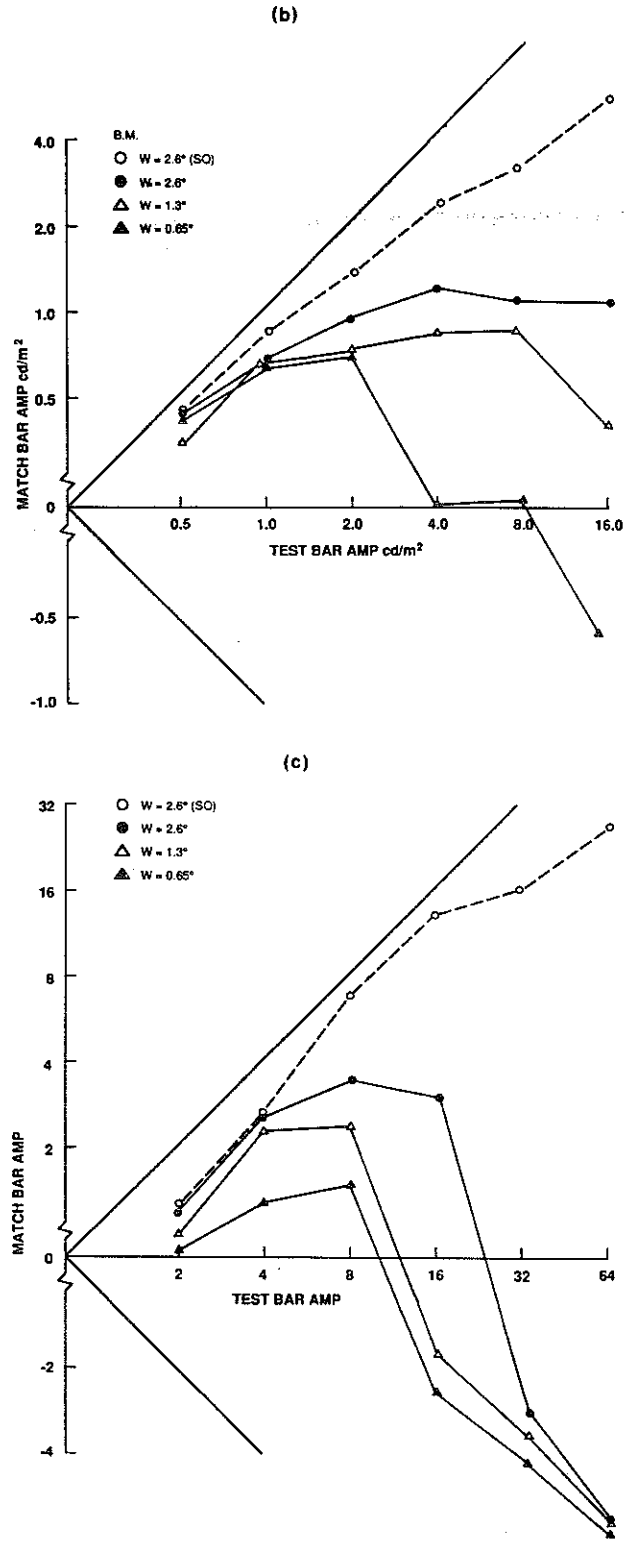

Figure 3. (*Contd.*)

(a)

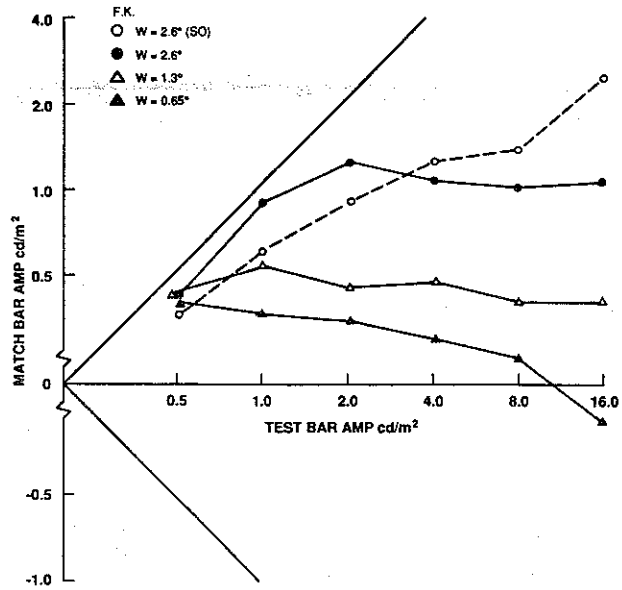

(b)

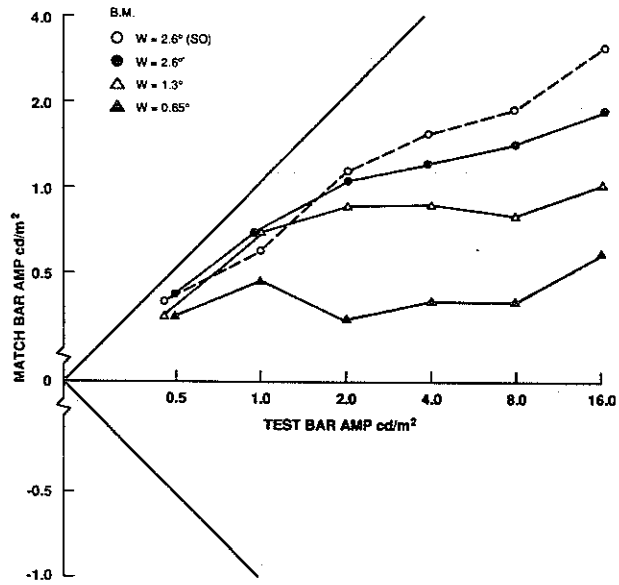

Figure 4. As in Fig. 3 except for stimuli with negative-going edges. Note that in this plot positive values of test-bar amplitude imply that the central test area appeared darkened, negative values that the test bar appeared lightened. (a) FK, (b) BM, (c) model predictions. The units in (c) are arbitrary.

(*Contd. on next page.*)

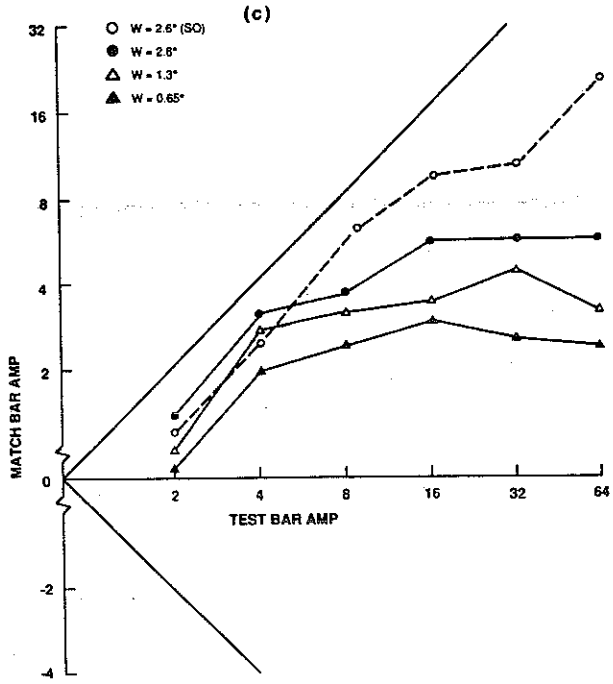

Figure 4. *(Contd.)*

negative gradient (IN) and the outer positive gradient (OP) stimuli are presented with the sign of their matched amplitude reversed in order that the shapes of their functions with respect to amplitude can be directly compared with the remaining two conditions. For the IN and OP conditions therefore, positive values show the degree of darkening of the test region, whereas for the IP (inner positive) and ON (outer negative) conditions, positive values show the degree of lightening of the test region.

DISCUSSION

The following are the main features of the data.

(1) For both inner gradient (IP and IN) stimuli an increase in gradient width results in a greater magnitude of induced brightness, whether the brightness induction was negative or positive.

(2) For the IP stimuli an increase in amplitude results in an increase in induced brightness followed by a downturn (except in the largest width condition for BM). In many instances the downturn goes through zero to negative implying that the central test region appeared darkened.

(3) For the IN stimuli there is no marked downturn at high amplitudes.

(4) For the OP stimuli the magnitude of darkening is greater at all amplitudes compared with the magnitude of lightening in the IP stimuli and moreover in the former there is no apparent downturn in the function at high amplitudes.

(5) There is little difference between the ON and IN conditions of the same gradient width.

We began with the intention of confirming previous reports of anisotropies in the magnitude of induced brightness in CCOB figures depending on (a) whether the stimulus had positive or negative going edges and (b) whether the sloping parts of the edge were on the inside or outside of the test region. We have confirmed that, at least at relatively high amplitudes, there is more induced brightness in negative-going than positive-going 'cusp' stimuli, in line with previous reports (Hamada, 1985). Moreover, we have shown that, at least with positive gradients, the magnitude of induced brightness is greater for outer compared with inner gradient stimuli. We now describe a model of brightness coding which is able to provide an account of the main results described above and to provide a quantitative fit to the data to a first degree of approximation.

A MODEL OF BRIGHTNESS CODING

The model we propose begins with the assumption that information about brightness is carried in the postconvolution response profiles of neural filters whose principal function is to detect the presence of local image features such as edges and bars. In the model we employ filters whose spatial weighting function is the second difference of a Gaussian (2DG):

$$F(x) = -A\left(\frac{x^2}{\sigma^2} - 1\right)\exp\left(\frac{-x^2}{2\sigma^2}\right), \tag{2}$$

where σ is the space constant and A the gain of the filter. The gain A was set to $1/\sigma$ throughout: this scaling factor was found to give the best fit to the data for the particular set of filters employed.

This weighting function is a one-dimensional approximation to that of the receptive field of a typical retinal ganglion cell. It must be pointed out at the outset that our use of such a filter does not imply any commitment to the idea that brightness computation is performed at the retinal level: indeed we have previously shown that most evidence points to the cortex as the site of brightness computation (Kingdom and Moulden, 1988), and a recent study by Shevell (1989) provides further persuasive evidence. Moreover, as we emphasise at the end of this communication, the physiological implementation of our model could well be via two classes of filter, namely bar (even symmetric) and edge (odd symmetric) filters in the cortex (Hubel and Wiesel, 1968; Burr *et al.*, 1989). We wish our model to be considered as a computational equivalent of processes at least some of which are almost certainly cortical in origin and which may involve at least two classes of feature detector.

In the illustrations given below, for clarity of exposition, we show the filter outputs as both positive- and negative-going; for example, Fig. 5 shows the familiar response of a 2DG filter to an edge and a bar: the edge response is biphasic and the bar response is triphasic. However, in keeping with the neurophysiological facts and with other computational theories, in our modelling we rectify the positive- and negative-going outputs of our model filter to reflect the fact that both 'on-centre' and 'off-centre' receptive fields respond to their preferred stimuli with an increase in firing rate; their output cannot, of course, fall below zero.

Up to this point there is nothing particularly novel in our approach. All current computational models involve a convolution of the input image with some sort of filter,

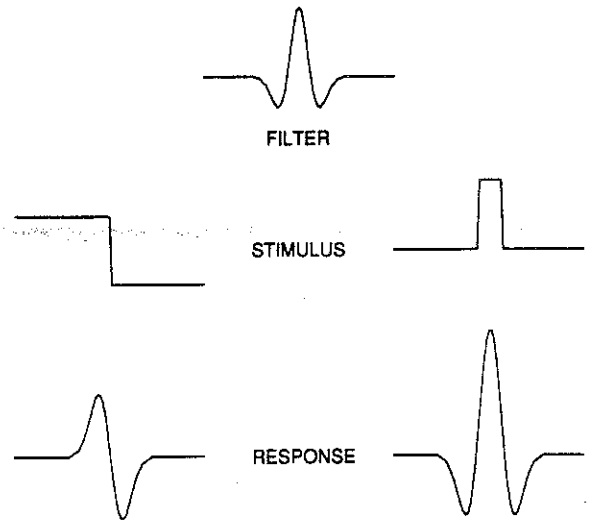

Figure 5. The response of a 2DG (second difference of a Gaussian) filter to an edge and a bar. The edge response is biphasic, the bar response triphasic.

and any even-symmetric filter will give a biphasic response to an edge and a triphasic response to a bar.

We hypothesize that in the neural image the correlate of the contrast of a bar is the magnitude of the central lobe in the triphasic response function, either its peak amplitude or its 'mass' (the area under the curve: its integral); we discuss these alternatives later.

In the case of an edge the correlate of contrast is the amplitude or mass of either of the two lobes in the biphasic response (the two are identical in the case of a balanced filter and perfectly correlated in the case of an unbalanced filter).

We also assume that there exist filters at different scales at any point on the retina and that their outputs are pooled in some way (this too is not a novel idea; the idea was first proposed by Wilson and Bergen (1979)); and further that the average size of the filters increases with retinal eccentricity.

The key novel features of our model are the following. First, we propose rules according to which the response profiles might be interpreted as representing either edge-like or bar-like stimuli. Second, we propose a 'default-rule' associated with triphasic responses which eliminates the need for a deconvolution operation on the overall filter response to recover the original luminance profile. Third, we incorporate a physiologically plausible thresholding operation based on signal-in-noise considerations. Fourth, the final description of brightness is computed as the average of the individual filter descriptions (unlike MIRAGE (Watt, 1988), in which filter outputs are pooled *before* the extraction of primitives). Finally, in order to account for certain brightness anisotropies, we assume that 'off-centre' filters are on average larger than their 'on-centre' counterparts.

The precise nature and significance of these features of our model will become clear in the following qualitative description of its operation, which we have expressed as a set of five 'rules'. Subsequently we give a quantitative description of its implementation

and of its fit to the data. In a later section we discuss its properties and strengths in comparison with those of a number of other current models of brightness coding.

Rule 1. Biphasic and triphasic response profiles have unique symbolic connotations
If the response profile of a filter is biphasic then it asserts the presence of an edge. If the response profile of a filter is triphasic then it asserts the presence of a bar.

Rule 2. Biphasic and triphasic responses carry different implications about stimulus brightness
(a) A biphasic response implies a step change in brightness whose magnitude is given by the peak or mass of its profile. (b) A triphasic response implies a localized change in brightness analogous to a rectangular function (though the precise shape may not be rectangular). The magnitude of the brightness increment (or decrement) is given by the peak or mass of the central lobe. In principle, information about the relative luminances of the regions on either side of the stimulus producing the triphasic response is also potentially available from a comparison of the magnitudes of the outer lobes. However, in our model this aspect of the triphasic response profile provides no information to the pooling mechanism that computes a final brightness value.

Rule 2b is the most extempore of our assumptions. We are led to propose it because we have discovered, despite an exhaustive inductive search for a computational rule, that the luminance profiles of our CCOB edges cannot be recovered by a deconvolution process involving the basic properties (peak, mass, standard deviation) of the three zero-bounded response lobes, even in a noiseless system. While the latter can give a good account for some classes of stimuli that generate a three lobed response (such as those with symmetry or with only sharp edges) we were forced to abandon any idea of such a deconvolution process in modelling our data, which are based on asymmetrical stimuli with one sharp and one graded edge.

Note that not all filter responses are biphasic or triphasic in their response to our stimuli. For example, Figure 9b shows that the largest filter gives a five-lobed response to an OP (outer-positive) stimulus. In our model the symbolic feature description of such a stimulus for this sized filter taken alone is that of a 'triple bar', since it consists of three overlapping triphasic responses. The resulting description of brightness by such a response is as shown in the figure. This brightness description is of course subsequently averaged with those of the smaller filters which do not produce a five-lobed response, and the 'triple-bar' characteristic disappears in the pooling process.

Rule 3. Small response lobes are thresholded out
Consider the response of a 2DG filter to a bar (Fig. 5). The output profile is a symmetrical triphasic function, the symmetry reflecting the fact that the two edges of the bar are identical. Compare this with its response to a CCOB stimulus, in which one edge is sharp and the other is graded (Fig. 6). The output profile is no longer symmetrical: the negative lobe corresponding to the graded edge is wider and—crucially—shallower than that corresponding to the sharp edge. The more gradual the edge, the more marked the effect. Clearly, if we apply a thresholding operation to these outputs at some point the shallow negative lobe will disappear, leaving only the large positive central lobe and the negative lobe corresponding to the sharp edge: the response is now biphasic, not triphasic, signalling 'edge', not 'bar'.

Rather than simply assuming some brutal, arbitrary, clipping procedure we see this

thresholding operation as an inevitable property of any noisy (and therefore any practical) signal-processing mechanism. A signal lobe has to be detected against a background of neural noise. As the amplitude of that lobe diminishes (with decreasing stimulus amplitude or increasing edge width) there will come a point at which it will simply be lost in the neural noise and effectively cease to exist: it has been thresholded out in our terms.

The removal of below-threshold regions of response activity is, according to our model, the mechanism by which illusory induced brightness in the CCOB stimuli occurs. It effectively loses some of the information about the presence of gradual luminance gradients, while preserving information about sharp ones.

Rule 4. Brightness descriptions from the individual filters are first independently computed and then pooled

Like others, we assume that filters exist at a range of different spatial scales. However, unlike MIRAGE (Watt, 1988) for example, in which filter responses are combined prior to the stage of symbolic description, in our model the filters act independently and the resulting symbolic description of brightness from each is computed separately and then averaged.

Rule 5. There are different spatial weighting functions for 'on-centre' and 'off-centre' filters

The data reported in this paper, and other experiments described earlier, clearly show anisotropies in the brightness response to positive and negative contrasts. In order to account for these anisotropies we suggest that the spatial weighting functions of 'off-centre' filters are wider than those for 'on-centre' filters. There is some support for this notion which we describe later.

Having outlined the key features of the model we will now illustrate its operation by means of a few examples; in the course of this exposition we shall discuss some aspects of the model in more detail.

Figure 6 shows the response of a 2DG filter to a CCOB edge of various gradient widths and amplitudes. As can be seen, as gradient width increases, the response becomes less bar-like (triphasic) and more edge-like (biphasic). The change from signalling 'bar' to signalling 'edge' is fairly sudden and occurs when the mass or peak amplitude of the critical lobe (the right-most one in the figures) falls below threshold, or, more specifically, gets lost in neural noise.

A given sized filter will thus signal 'edge' to a wide-gradient CCOB edge, but 'bar' to a narrow-gradient CCOB edge, and this is the model feature that enables us to account for the effect of gradient width on the magnitude of induced brightness seen with the positive (IP) and negative (IN) 'cusp' stimuli.

Figure 6 also shows how increasing amplitude produces the opposite change from that produced by increasing gradient width: as amplitude increases a point is reached when the third critical lobe goes above threshold and the signal changes from 'edge' to 'bar'. The effect of the change from 'edge' to 'bar' at a critical stimulus amplitude is the feature which allows us to account for the collapse of the illusion at high amplitudes which occurs with our OP (outer positive) stimuli.

A single-filter model based upon the above rules predicts the overall qualitative effects of both gradient width and amplitude shown in Fig. 3a,b and Fig. 4a,b. However, the computed functions do not reflect the gradual change in the empirical

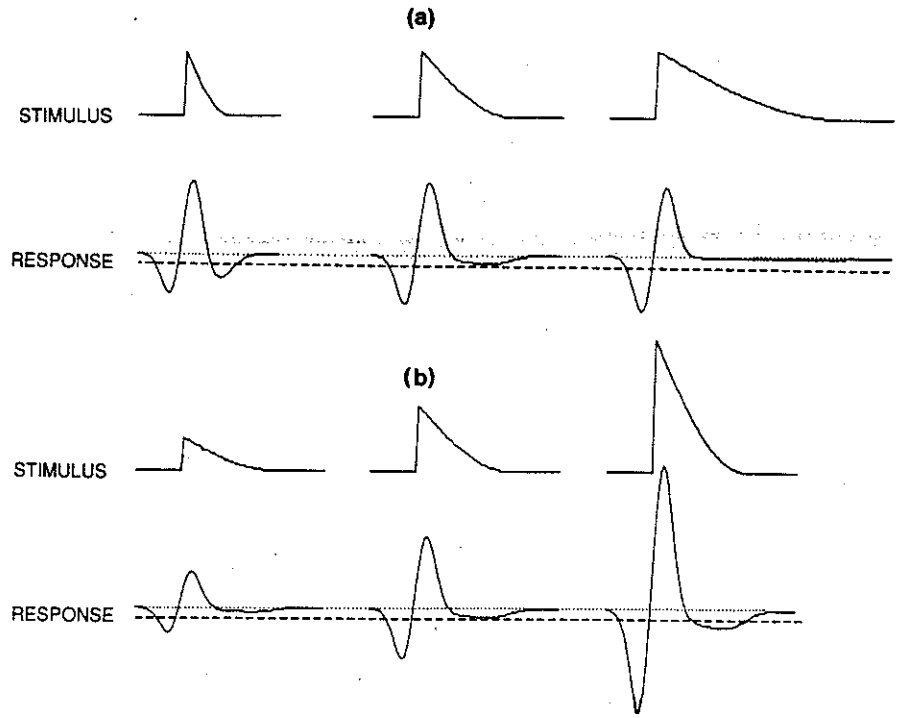

Figure 6. The response of a 2DG filter to CCOB edges varying in (a) edge width and (b) amplitude. The zero response level is indicated by the dotted line while the threshold level on the negative lobes is indicated by a dashed line. Only the smallest edge width condition in (a) produces a triphasic 'bar' response, the remaining edge width conditions producing an 'edge' response. In (b) only the largest amplitude condition produces a triphasic 'bar' response, the remaining conditions produce a biphasic 'edge' response.

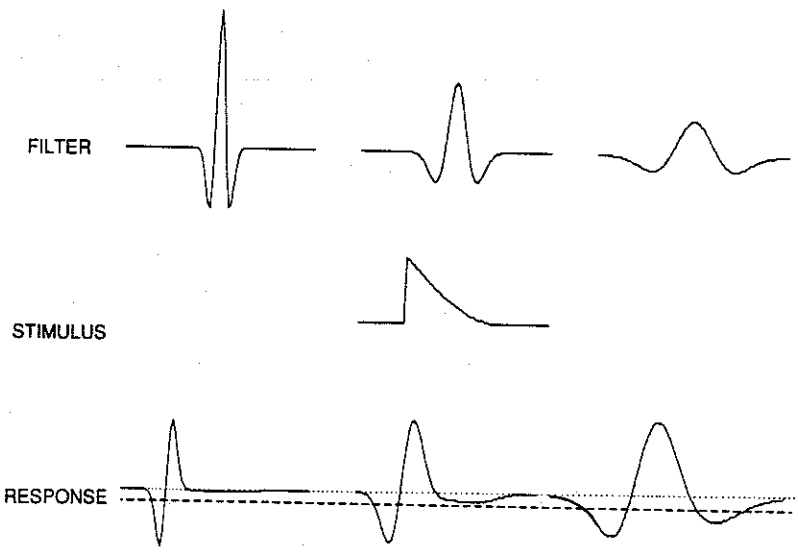

Figure 7. Effect of filter size on feature encoding with a CCOB edge. Because of thresholding, only the largest filter produces a triphasic 'bar' response. The two smaller filters produce biphasic 'edge' responses.

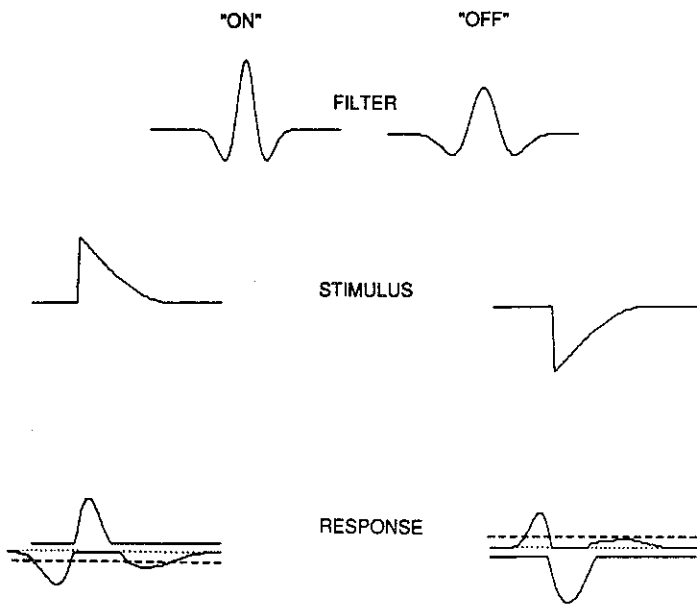

Figure 8. Different sizes of 'on' and 'off' centre filters result in a different feature code for a positive as opposed to negative gradient CCOB stimulus. Note that the spatial weighting function of the 'off' centre filter is illustrated as if it were an 'on' centre unit. This is in order that its rectified response profile can be shown as a negative response, below that of the 'on' centre filter. After thresholding, only the left hand positive gradient CCOB stimulus produces a triphasic 'bar' response, while the negative gradient CCOB produces a biphasic 'bar' response.

functions that occur as stimulus amplitude changes. The computed functions predict a much sharper increase in the magnitude of induced brightness with stimulus amplitude, followed by a sudden and dramatic decrease. This mismatch disappears if we suppose that a range of filters of different spatial scale is involved.

Figure 7 shows the response of a set of different-sized filters to a given CCOB edge. As can be seen, while the two smallest filters will signal 'edge', the largest will signal 'bar'. If one were to increase the gradient width, the large filters would one after another change their signals to 'edge'. If one were to increase stimulus amplitude, the smaller filters would one after another change their response to 'bar'. This feature of the model, whereby a given stimulus produces different symbolic descriptions in the different-sized filters, together with the hypothesis that the final description of brightness is computed as the average of those individual filter descriptions, are the properties that emulate the smoothly changing functions seen in the data, particularly in the wide-gradient conditions of Fig. 3 and 4.

The anisotropies between the positive and negative stimuli occur, we suggest, because the receptive fields of 'on' and 'off' centre filters have different spatial weighting functions. In keeping with electrophysiological evidence (Schiller, 1986) we assume that the filter responses are rectified in the sense that the positive and negative components of the filter response are carried by increases in the firing rate of 'on' and 'off' centre filters respectively. We suggest that the 'off' centre filters have on average larger space constants. The effect of this anisotropy on the processing of the stimuli employed here is illustrated in Fig. 8. In this figure, the positive CCOB stimulus is encoded as a 'bar'

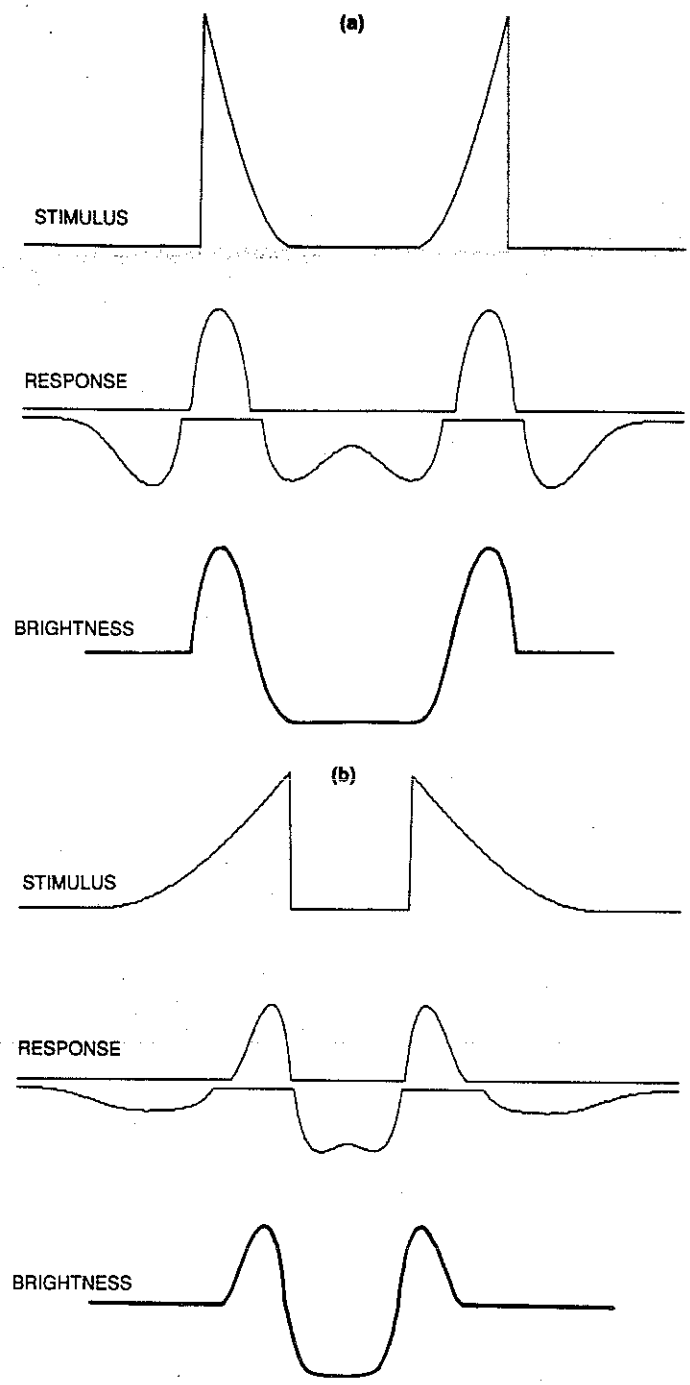

Figure 9. The symbolic description of brightness from the largest filter for: (a) a high amplitude CCOB stimulus with a positive inner gradient; and (b) a CCOB stimulus with a positive outer gradient. In both, the presence of a central negative response lobe is interpreted as a negative going central bar, whose brightness description is then averaged with those from the smaller filters. Because of the contribution from the large filters, the central test region appears darkened in both (a) and (b).

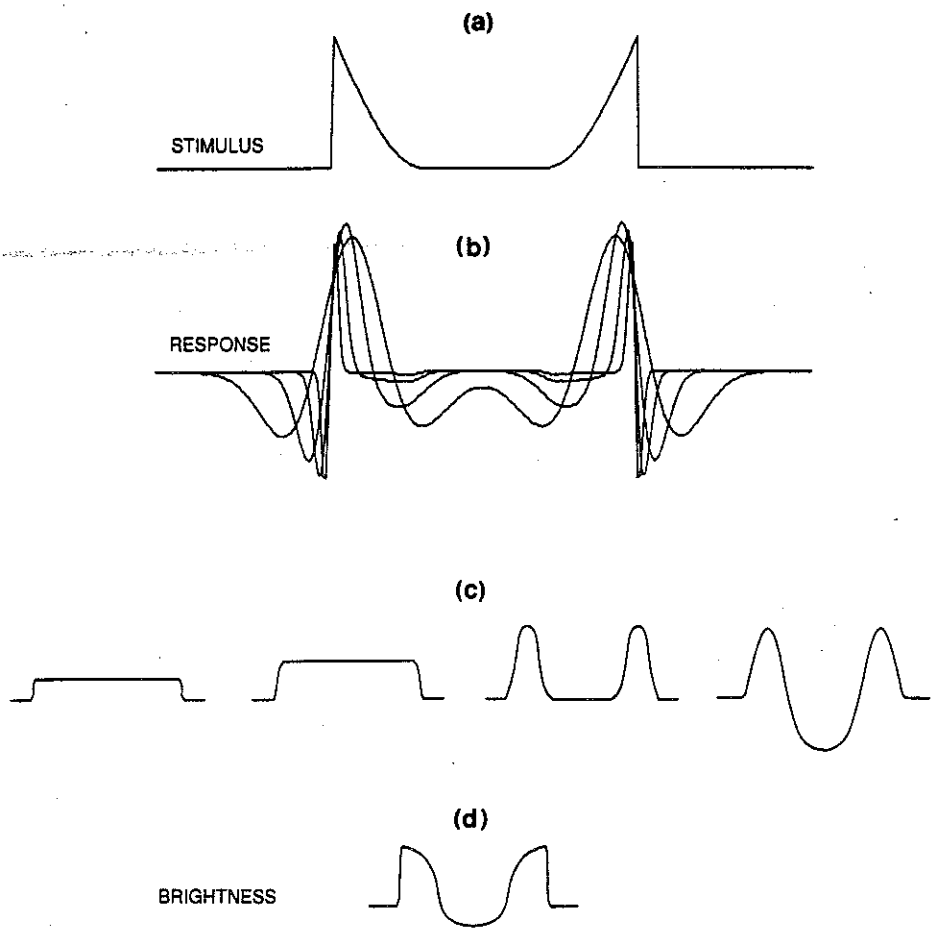

Figure 10. The individual responses (b) to a CCOB stimulus (a) with a positive inner gradient for the range of filters employed. (c) illustrates separately the symbolic descriptions of brightness produced from each filter response while (d) shows the final computed brightness description, which is the average of the descriptions in (c).

because the 'off' centre filter response produces two above-threshold lobes, thus producing a triphasic response for the stimulus as a whole. On the other hand, for the negative CCOB stimulus, the smaller 'on' centre filter produces only one above-threshold lobe and hence signals an 'edge'.

Figure 9 illustrates how the model accounts for: (a) the change from induced lightening to induced darkening at high contrasts in the PI stimulus; and (b) the lack of a turnover at high contrasts for the stimuli with outer gradients. In both cases the largest sized filter signals the presence of a central bar whose contrast in Fig. 9a is in the opposite direction to that of the induced lightening while in Fig. 9b it is in the same direction as that of the induced darkening.

Figure 10 illustrates how the brightness descriptions for the different sized filters are combined to produce the final brightness percept.

For the purposes of this paper we have not attempted to describe quantitatively the

complete brightness profiles of our CCOB stimuli as encoded for each filter output, nor therefore the final brightness profile resulting from their combination. To do this would require at the very least a description of the degree of blur of the features (edges and bars) encoded by each filter. This is outside of the scope of our model as it currently stands: its specific requirement was to describe quantitatively only the difference in brightness between the central test region and the background. The symbolic brightness profiles shown in Figs 9 and 10 are therefore meant to be only approximate descriptions, except in the region of the background and central test areas. Note also in Fig. 10c that the brightness profile encoded by the two smallest filters is homogenous for the central test region. For these filters the discontinuities in the stimulus produced biphasic responses and were therefore encoded as edges producing step changes in brightness. Since there were no discontinuities between those edges, the step changes were continued unabated across the central test region. This can be regarded as a 'filling-in' process, but only at a symbolic level. It is a computational procedure, and does not imply any spreading of neural excitation between the edges to produce a pattern of neural activity isomorphic with the brightness profile (see section entitled 'Comparison with other models' for a further discussion of the nature of 'filling in' processes).

Details of model implementation

To implement our model we employ a thresholding operation which simulates a plausible physiological mechanism, and is essentially identical to the thresholding mechanism employed in the MIRAGE model of Watt and Morgan (1985). Random noise is added to the filter response prior to rectification: the response at each point along the response axis $R(x)$ is allowed to deviate by a random amount within the limit imposed by the noise level Q such that $-Q < R(x) < Q$. The response is then separated into its above-zero and below-zero components $R(x)$ and $-R(x)$, and potential signals in both arrays are identified as areas of response activity that are bounded by zero response on either side.

These potential signals have both a mass (area) and a peak amplitude. Whichever of these primitives is chosen to encode brightness, all signals whose mass or peak amplitude falls below a threshold level are ignored. The threshold level was chosen to be that level which just ensured that signals could not occur by chance ($p < 0.01$) in noise alone. The model predictions shown in Figs 3c and 4c employed mass as the code for brightness because the fit was superior to that of peak amplitude for the particular filter parameters chosen. Nevertheless, a model assuming peak amplitude as the primitive can also simulate the overall differences found between the different classes of stimuli we used and at this stage we are not able to rule out peak amplitude as the primitive for brightness. However, our initial findings indicate that in order to make peak amplitude provide a good fit to the data the filters have to be significantly larger than the ones required if mass is assumed to be the primitive. See below for a discussion of the filter sizes we employed.

Figure 11 illustrates the processes of adding noise and signal extraction. We applied a compressive non-linearity on the *output* to the filter, in keeping with neurophysiological evidence (Robson, 1975) as well as psychophysical evidence showing Weber-like behaviour for contrast discrimination thresholds (Wilson, 1980; Whittle, 1986). We have ignored the effects of the non-linearities in photoreceptor transduction which result from retinal adaptation to the prevailing light level (Shapley and Enroth-

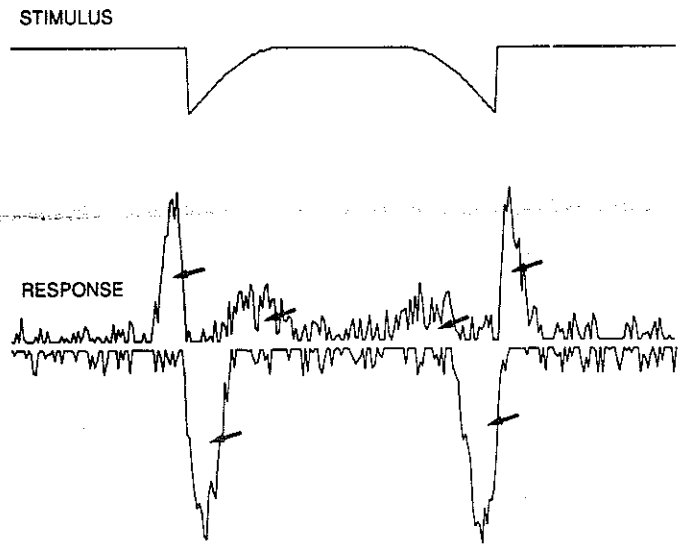

Figure 11. The mechanism of thresholding. Random noise is added to the filter response, and areas of zero bounded activity in both the positive and negative profiles are isolated. Only those whose peak amplitude or mass exceeds a threshold are then accepted as signals, and these are indicated by arrows. This mechanism of thresholding follows that of MIRAGE (Watt and Morgan, 1985).

Cugell, 1984) since mean background luminance was fixed throughout the experiment. The output non-linearity is achieved by transforming the response $R(x)$ logarithmically using the equation

$$R'(x) = \text{sgn}(R(x))\ (1 + \ln(|R(x)|)). \tag{3}$$

The effect of such a transform allows slightly smaller filters to be used in the model, since the transform has most of its effect in reducing the magnitude of the very tall lobes signalling 'edge' in the small filters.

Estimates of test-background differences were also made for the match stimulus over a range of amplitudes. Since both test and match stimuli were on the same background, the amplitude of the match stimulus which was computed to give the same test-background brightness difference as that of the particular CCOB condition under investigation is taken as the model prediction for that condition.

The model predictions of test-background brightness difference for a given stimulus were computed by averaging the estimates over 100 trials, with 25 iterations for each of the four different-sized filters. On each trial fresh noise was added to the filter response.

Range of filter sizes employed

For the stimuli with outer gradients, whose sharp edges were the closest to the fovea, the space constants (see equ. 2) for the 'on' centre cells were 2, 4, 8, and 16 arcmin; for the 'off' cells, they were 3, 6, 12 and 24 arcmin. These values resulted in excitatory centre widths which ranged from 4.7 to 37 arcmin for the 'on' filters, and inhibitory centre widths of 7.1–56 arcmin for the 'off' filters. For the remaining stimuli the filter sizes were scaled upwards depending on the position of the sharp discontinuity of the stimulus in peripheral vision (central fixation was employed throughout the experiments). On the basis of the rule of thumb that filter sizes roughly double with every 4 deg of eccentricity

(Wilson, 1978), the scaling resulted in space constants for the widest gradient condition in the range 3.1–25.0 arcmin for the 'on' centre filters and 4.6–37 arcmin for the 'off' centre filters. This method of scaling may have exaggerated the effects of eccentricity since the gradient part of the stimulus, being on the inside, would stimulate slightly smaller filters than the ones at the sharp discontinuity. Nevertheless, the scaling gives an engineering approximation to the changes in filter size that one might expect for the processing of the stimulus as a whole.

Model predictions

The predictions of the model are shown in Figs 3c and 4c, alongside the empirical results. The fit is closer to the data of FK than those of BM, though a better fit to BM's results could easily be obtained by reducing the size of the 'off' centre filters by a constant. It is, however, more likely that the differences between subjects reflect criterion differences rather than differences in the properties of the underlying mechanisms. We have not attempted to finely tune the parameters of the model to fit the data exactly; our aim is simply to show that the principal qualitative relationships between the various conditions are accounted for by the model.

Other evidence for the model

The model of brightness that we have outlined above is of necessity speculative, and there is little external evidence for some of its assumptions. Nevertheless, there are a couple of previous findings which are consistent with the main features of the model.

Firstly, Tolhurst (1972) suggested that there might be reciprocal inhibition between edge and bar detectors, in order, for example, to prevent bars being seen on either side of an edge (both 'bar' and 'edge' detectors are stimulated by an edge). He also suggested that such reciprocal inhibition might underlie the observation that a thin line positioned along an edge substantially reduces the apparent contrast of that edge. This finding is exactly what would be predicted by the model described here. The thin contour would produce a 'bar' signal in the smallest filter, and therefore remove its contribution to the perceived contrast of the edge. However, some perceived contrast would remain since the larger filters would not be stimulated by the thin contour, but would be driven by the edge.

Secondly, there is some independent evidence in support of the suggestion that 'off' and 'on' centre cells may have different spatial weighting functions. Du Buf and Roufs (1985) measured detection thresholds and brightness for both incremental and decremental discs as a function of disc area against a fixed reference background. The area beyond which there was no further effect of area on brightness was found to be about 0.6 log area units greater for the decrements than for the increments (in our modelling the 'off' filters were about 0.4 log area units greater than the 'on' units). Du Buf and Roufs suggested that this indicated a broader point spread function for decrements. There is also some supportive neurophysiological evidence for larger 'off' centre receptive fields (Hammond, 1974; Van Erning *et al.*, 1988).

Anisotropies in the brightness and saliency of stimuli depending on whether they are greater or lesser in luminance than the surround have been frequently reported in a number of studies of brightness phenomena: Heinemann (1972) and Moulden and Kingdom (1989) with simultaneous contrast; Jory and Day (1979) with Kanizsa's triangle; Spillman and Levine (1971) with the Hermann Grid illusion; Legge and Kersten (1983) with contrast discrimination; Kingdom and Moulden (1989) for the

detection of line signals in visual noise, to cite but a few examples. The anisotropy has been attributed to a nonlinear, Weber-like transform of input intensity at an early stage in visual processing, which, if taken into account, gives symmetry to the results of some of the data with increments and decrements (Legge and Kersten, 1983; Burkhart *et al.*, 1984; Whittle, 1986; Kingdom and Moulden, 1989). Essentially, the difference in the log transformed luminance between the signal and its background (the suggested correlate of perceived contrast) will have a larger absolute value in the case of a decrement than for an equal amplitude increment. However, such an explanation could not account for the anisotropies between the CCOB figures with positive and negative gradients found here and elsewhere (nor indeed will it for some of the other anisotropies mentioned above as we have discussed before, Moulden and Kingdom, 1989). First, consider the range of amplitudes within which the CCOB illusion is manifest: if the negative cusps were represented internally with a higher value than their equal-amplitude positive cusp counterparts, so also would their respective match stimuli, and thus the anisotropy would not be expected to reveal itself. Second, consider the stimulus amplitude at which the illusory brightness would be expected to break down on the nonlinear transform interpretation. If the internal representation of the negative cusps was of a higher value than the positive cusps, one would expect the illusion, if anything, to collapse at a *lower* amplitude than for the positive cusps, the opposite of what is found.

We turn now to a comparison of our model with other current models of brightness coding. We (Kingdom and Moulden, 1988) have previously categorised brightness models into essentially three classes: contrast sensitivity function (CSF) models, lightness integration models, and edge detector models; the model described here falls into the third category, which we shall refer to below by the more general term 'feature detector' models. We will discuss the other candidate models within this framework.

CSF (Contrast Sensitivity Function) models

It is axiomatic that low-frequency attenuation in contrast sensitivity is ultimately responsible for the illusory induced brightness in CCOB figures. However, the question is: can the CSF, which contains an explicit description of such attenuation, be used to model the magnitude of induced brightness quantitatively as a function of, for example, contrast? We have previously reviewed the extent to which the CSF can successfully predict the pattern of induced brightness in CCOB figures (Kingdom and Moulden, 1988), and so only a few comments will be made here. Essentially, the CSF is reasonably accurate in predicting the appearance of CCOB figures such as the MF (missing fundamental) square-wave (Campbell *et al.*, 1971) and the Cornsweet figure (Cornsweet, 1970) up to contrasts at which the figures become distinguishable from their step-edge relatives (in the case of an MF stimulus, a normal square-wave; in the case of a Cornsweet stimulus, a step-edge).

A simple and plausible rule suffices to account for the appearance of low contrast CCOB figures. This is the 'default to square-wave' rule and it states that, in the case of the MF display for example, the stimulus will appear indistinguishable from a square-wave until its missing component (the fundamental) would become independently detectable (Campbell *et al.*, 1971). The problem with the CSF approach, however, is that at contrasts above that at which the missing low-frequency components become independently detectable in CCOB figures in general (and therefore at contrasts at which such stimuli are distinguishable from their step-edge relatives) illusory

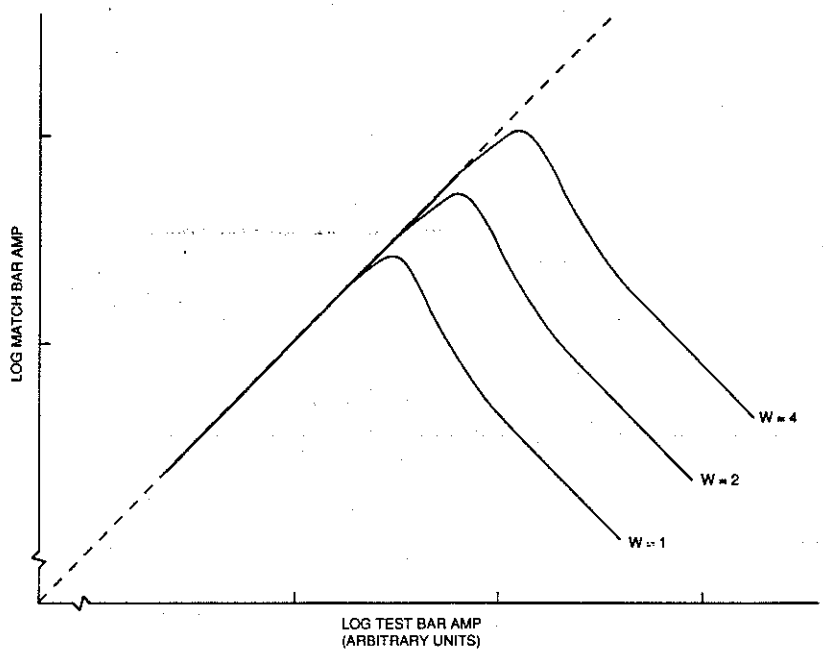

Figure 12. Predictions from the Retinex model of Land and McCann (1971). The graph plots the (log) match bar amplitude which matched a CCOB stimulus for computed brightness as a function of the (log) amplitude of the latter, for three gradient widths (*W*). The shapes of the functions would be the same (except for polarity) for all the four CCOB stimuli employed here.

brightness continues to exist, and the simple 'default to square-wave rule', by definition, becomes inapplicable. To overcome this problem, Dooley and Greenfield (1977) employed a modification of the rule in which the difference in brightness between a Cornsweet figure and its step-edge relative was correlated with the difference in their respective power spectra (after attenuation by the CSF). This, however, now produced an overestimation of the magnitude of induced brightness at high contrasts: a reasonable description of the data was only possible using a CSF which changed its shape with contrast, specifically by reducing the degree of low-frequency attenuation monotonically with increasing contrast. While this latter manipulation is consistent with some psychophysical evidence (Blakemore *et al.*, 1973; Georgeson and Sullivan, 1977, though see Kulikowski, 1976) it is clear that the initially most attractive feature of the CSF approach (its simplicity) has now been lost. Leaving that aside, the fit to the data that Dooley and Greenfield achieved still has only the status of a mathematical description, rather than a physiologically plausible computational model. The key issue remains unresolved: how might the visual system compute the properties of the (transformed) power spectra that would allow the relevant comparison of the CCOB stimuli with their step-edge relatives? Given the (now almost universally accepted) view that the description of the spatial frequency content of the visual scene is carried in a piece-wise fashion by spatially localised band-pass filters in the cortex, rather than as a global description, it is clear that the CSF approach to CCOB illusory brightness has little predictive value in terms of understanding the nature of the underlying mechanisms (other than to say that they are linear at, or near to, threshold).

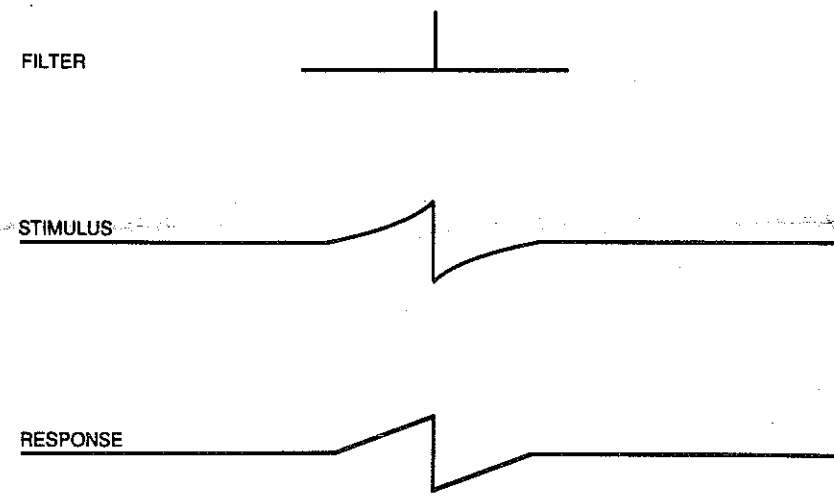

Figure 13. Predictions for the Land (1985) brightness model. The stimulus is convolved with a balanced filter consisting of a one pixel wide excitatory centre and a 40 pixel wide inhibitory surround. Brightness is encoded directly as the response of the filter. The CCOB illusion is not predicted for a Cornsweet (1971) figure.

Lightness integration models

This important class of brightness models originated from the need to account for 'lightness constancy': the apparent constancy of lightness and hue in the context of wide variations in the intensive and spectral content of illumination. All lightness integration models employ four stages. They are: (i) a nonlinear (usually log) transform of the luminance profile; (ii) differentiation of the (transformed) luminance profile; (iii) thresholding; and (iv) reintegration. The key property of the models in the present context is the third, thresholding, stage. This stage has the purpose of removing any gradual luminance gradients in the image, such as, for example, gradual gradients in illumination. Thus in principle lightness integration models will predict the Cornsweet illusion since the thresholding stage will remove the signal for the gradient part of the edge, while preserving the sharp discontinuity.

Figure 12 shows the predictions generated by an implementation of the earliest lightness integration model, the Retinex model of Land and McCann (1971), for the stimuli employed in the experiments described here. While the effects of gradient width and amplitude are qualitatively predicted for the inner gradient stimuli, the shapes of the functions are clearly inaccurate. More importantly, however, the Retinex model would predict no difference in the magnitude of brightness for the inner- and outer-gradient conditions. The more sophisticated versions of lightness integration models (Horn, 1974; Blake, 1985) will also suffer from this defect.

Land's (1985) model

A common criticism of lightness integration models is their failure to predict simultaneous contrast effects (Hurlbert, 1986; Reid and Shapley, 1987). Indeed the aim of such models is the correct recovery of reflectance, whereas the phenomenon of simultaneous contrast implies errors in reflectance recovery. Partly to overcome this criticism, Land (1985) has suggested that brightness may be modelled using a filter with

a small excitatory centre and large inhibitory surround. The brightness of the image is given as the logarithm of the convolution output of the filter with the input luminance profile: no deconvolution or integration process is required. Strictly speaking, therefore, Land's (1985) model is not a lightness integration model. We have included it in this section, however, since it developed out of the Retinex model. In the model gradual illumination gradients are eliminated by virtue of the fact that the filter is unresponsive to linear gradients in luminance. Simultaneous contrast effects are predicted. However, the Cornsweet illusion is not predicted by this model as is illustrated in Fig. 13.

Feature detector models

Like the model we have proposed here, all feature detector models which deal with brightness computation begin with the idea that brightness coding is linked to the mechanism which generates a symbolic description of salient local features in the image such as edges and contours. The feature detection model which we examine first is the one which has dealt most explicitly with brightness: the neural network model of Grossberg and his co-workers (Cohen and Grossberg, 1984; Grossberg and Todorovic, 1988).

Grossberg's spatial averaging model

The central feature of this model, and of an essentially similar model by Hamada (1984), is that brightness is encoded by the average response of an even-symmetric filter within regions of the image bordered by sharp luminance discontinuities. Since the correlate of brightness is this average response we refer to the model as a spatial averaging model. Grossberg and his co-workers place much emphasis on the presumed neurophysiological implementation of the spatial averaging process: in their model it occurs as a result of a diffusive spread, or 'filling-in', of neural activity within a contour-bounded region, resulting in a region of homogenous neural activity which isomorphically defines the brightness profile on a point by point basis. We (Kingdom and Moulden, 1988) and others (Foster, 1983; Laming, 1983) have argued against the logical requirement of such a filling-in process, whose proponents implicitly assume that it is necessary to produce an isomorphic spatial mapping of internal response to percept. However, the suggestion that the spatial average is the code for brightness is not in fact contingent on the existence of a filling-in mechanism such as that postulated by Grossberg and his co-workers, and therefore remains a viable possibility.

A necessary requirement of spatial averaging models is that the filter whose output is averaged is unbalanced. Such models employ a filter whose spatial weighting function is an approximation to the receptive field of a retinal ganglion cell, but whose excitatory centre is more heavily weighted than the surround. The filter therefore gives a DC response to homogenous light stimulation. This is necessary because if the filter were balanced the brightness response for a homogenous disc computed by spatial averaging would rapidly approach zero as the disc increased in size.

Grossberg's spatial averaging model qualitatively predicts the Cornsweet illusion because the part of the model which detects the contours within which spatial averaging takes place is thresholded. Only the sharp discontinuity is reflected as a contour, the gradients on either side falling below threshold. The spatial averaging process results in a difference in average response activity in the regions on either side of the sharp Cornsweet edge, thus producing the illusory brightness difference.

This model has the important advantage over lightness integration models in that it predicts some important phenomena such as simultaneous brightness contrast. Unfortunately, only qualitative demonstrations have been used to test the model, and no quantitative data on CCOB stimuli have been produced. Since the exact implementation of the model is complex and indeterminate we have not attempted to apply it to our results. The data reported here would present an informative challenge to the model; it would be interesting to see how it might meet that challenge.

While it is not the purpose of this report to provide a detailed critique of the spatial averaging model, it is worth noting a couple of the more important criticisms. Firstly, although most retinal ganglion cells are slightly unbalanced, resulting in a response to even illumination, this is simply not the case for cortical cells which invariably do not respond to homogenous light stimulation (Robson, 1975): and, as we have said, it is almost certainly the cortex which is the site of brightness computation. Secondly, the model cannot account for 'transitivity' effects in induced brightness. Arend *et al.* (1971) first showed that the induced brightness on either side of a radial Cornsweet edge 'carried acrosss' suitably placed rings in those regions, as well as into the rings themselves (see also Kingdom and Moulden, 1988). Since the spatial averaging process in Grossberg's model is contour-bounded, such an effect would not be predicted.

Kingdom and Moulden (1988) model for the Cornsweet illusion

In this model the brightness difference across a Cornsweet edge was given by integrating the masses of zero-bounded regions in the response profile of a 1DG (1st Difference of a Gaussian) filter. The 1DG model, while essentially in the 'feature detector' tradition, had some features in common with the Retinex model: the employment of a 1DG filter and the integration of zero-bounded response regions are formally identical to the differentiation and integration stages in the Retinex model; the principal difference between the two models lay in the nature of the thresholding stage.

As in the modified model proposed here to account for our more recent data, the illusory brightness in the Cornsweet figure was held to result from the curtailing of lobes in the filter response by noise. While the original model provided a good fit to earlier data collected on the Cornsweet illusion (see Kingdom and Moulden, 1988) it cannot account for the anisotropies of induced brightness in the CCOB stimuli described here. Specifically, the simple integration of zero-bounded masses from a 1DG (or indeed any) filter cannot predict simultaneous contrast effects, and it is difficult to see how any differences between positive and negative-going stimuli could be produced by any odd-symmetric filter.

Watt and Morgan's MIRAGE model

Although this model (Watt and Morgan, 1985) was formulated principally to account for data obtained with positional acuity tasks, Watt (1988) has provided a hint of how brightness might be extracted from the symbolic description that the MIRAGE algorithm generates. In MIRAGE, the positive and negative components of the responses of an array of 2DG (2nd Difference of a Gaussian) filters of different space constants are separately combined prior to the stage at which zero-bounded response regions are isolated. From such regions a symbolic feature description is generated, with biphasic and triphasic responses being encoded as edges and bars respectively, as in our model. Watt (1988, p. 42) suggests that the amplitude of luminance change (and hence its correlate in brightness) of bars and edges might be signalled by the product of

the mass and the standard deviation of the relevant response lobes (for a bar this would be the central lobe). However it is not clear what contribution, if any, the two outer lobes of a triphasic (bar) response make to the computation of brightness. As the model stands, these would have to be taken into account in some way, otherwise any luminance difference on either side of a bar could never be recovered during the brightness computation. However, as we noted earlier, we have been unable to discover any computational rule for recovering the luminance profile of stimuli with asymmetric edges (like CCOB edges) from the properties of the three zero-bounded response lobes in the outputs of noise-free 2DG filters.

In our model we avoid the need for recovering brightness information from all three lobes in a triphasic response. We do this by allowing a description of brightness to be computed separately for each size of filter. The luminance difference on either side of a thin contour will be carried by the larger filters which will be unresponsive to the contour itself. There will be, as found perceptually, a reduction in brightness contrast across the contour since the small filters will signal 'bar' (without the contour they would signal 'edge' and add their brightness contrast signal to the pooled response). It is hard to see how MIRAGE would account for the anisotropies in the CCOB stimuli described here without the incorporation of additional rules similar in effect to our pooling rules. What we conclude is not that MIRAGE is in this respect incorrect, but merely that it is incomplete.

Morrone and Burr model of phase congruence

This model (Morrone and Burr, 1988) offers an elegant and plausible explanation of how the visual system might detect both the presence and the location of salient features in the image, particularly edges and bars. The starting point of the model is that the position of edges and bars in the image can be determined from the points at which their individual Fourier components come into phase alignment with each other. The idea is that these points of 'phase congruence' are also detected by the human visual system. They are detected as peaks in 'local energy', where local energy is defined as the square root of the sums of squares of the responses of both odd and even-symmetric filters in the cortex. The nature of the feature is obtained by comparing the position of the peak in local energy with the position of the peak response from each of the two classes of filter. If the peak in local energy coincides with that of the even-symmetric filter, the stimulus is a bar; if it coincides with the peak of the odd-symmetric filter, the stimulus is an edge. The model has the particular advantage over other feature detection models in that it agrees well with both neurophysiological and psychophysical evidence for edge and bar detectors in the visual cortex (Hubel and Wiesel, 1968; Burr *et al.*, 1989).

However, it is by no means clear how it might account for the anisotropies we describe here or indeed for the appearance of Cornsweet edges in general. The critical test which we argue would defeat the model is the effect of contrast on the appearance of a Cornsweet edge. As a number of investigators have shown (for example Isono, 1979; Burr, 1987) the magnitude of induced contrast in a Cornsweet edge as a function of its amplitude follows an inverse U-shaped function. Moreover, the appearance of the actual discontinuity in the Cornsweet changes with amplitude: at low contrasts it is decidedly edge like, while at high contrasts it appears as a double (a bright and a dark) bar. Since the phase-congruence model is linear, the positions of the peaks in both the odd- and even-symmetric filter functions, as well as the local energy function, will

remain invariant with stimulus amplitude. Moreover, the relative amplitudes of those peaks will remain in strict proportion. If both the positions and relative amplitudes of all three signals remain invariant with stimulus amplitude, the feature content of the stimulus should also not change. It is also worth noting in passing that the phase-congruence model would make the rather implausible prediction that sinusoidal gratings are featureless since $(\sin^2 + \cos^2) = 1$ everywhere (M. A. Georgeson, pers. comm.).

SUMMARY AND CONCLUSIONS

We conclude that current models of spatial vision that address the issue of brightness coding (including our own earlier model) are unable to account for the quantitative data we have gathered concerning the effects of gradient width and amplitude on the magnitude of induced brightness in stimuli with CCOB edges.

We are led to propose an extension of our original model. The important new features are first that the final brightness signal is computed by pooling the outputs of filters at different spatial scales whose individual brightness signals are independently computed *before* pooling, and second that triphasic ('bar') signals carry information only about the contrast of the bar and not information about any differences that may exist between the luminances on either side of it.

This model has all the explanatory power of our earlier model, and is in addition able to give a good quantitative account for the data reported here. We need to modify a previous conclusion (Moulden and Kingdom, 1989) that computational models of brightness coding must be prepared to integrate the outputs of more than one class of spatial filter: we now conclude that biphasic and triphasic responses are pooled according to different rules.

Two final points need be made. First, while the model employs the same filter to signal the presence of both edges and bars, it is quite probable that at the physiological level there are separate edge and bar detectors, as we stated earlier. If this is so, then one of a number of possible ways that our model could be implemented physiologically would be through inhibition of edge detectors by same-scale bar detectors. This would implement the rule we employed which states that the presence of a triphasic response in the 2DG filter of given spatial scale implies no brightness difference across the discontinuity.

Second, it should be noted that the brightness coding rules in the model that we have described are essentially *post hoc*. Although precise and computational, the model does not itself constitute part of a computational theory, in that the rules we propose derive from our attempt to model one particular set of data rather than flowing from a full theory of brightness coding. However, as we argued in the text, the model has a number of strengths; not the least of these is its explicitness and consequent susceptibility to empirical evaluation.

Acknowledgements

This work was supported by the Science and Engineering Research Council (SERC) grant under the auspices of the Special Initiative on Image Interpretation (GR/D 89165).

REFERENCES

Arend, L. E., Buehler, J. N. and Lockhead, G. R. (1971). Difference information in brightness perception. *Percept. Psychophys.* **9**, 367–370.

Blake, A. (1985). Boundary conditions for lightness computation in a Mondrian world. *Comp. Vision Graph. Image Process.* **32**, 314–327.

Blakemore, C., Muncey, J. and Ridley, R. (1973). Stimulus specificity in the human visual system. *J. Opt. Soc. Am. A* **58**, 1300–1309.

Burkhardt, D. A., Gottesman, J., Kersten, D. and Legge, G. E. (1984). Symmetry and constancy in the perception of negative and positive luminance contrast. *J. Opt. Soc. Am. A* **1**, 309–316.

Burr, D. C. (1987). Implications of the Craik–O' Brien illusion for brightness perception. *Vision Res.* **27**, 1903–1913.

Burr, D. C., Morrone, M. C. and Spinelli, D. (1989). Evidence for edge and bar detectors in human vision. *Vision Res.* **29**, 419.

Campbell, F. W., Howell, E. R. and Johnstone, J. R. (1978). A comparison of threshold and suprathreshold appearance of gratings with components in the low and high spatial frequency range. *J. Physiol.* **284**, 193–201.

Campbell, F. W., Howell, E. R. and Robson, J. G. (1971). The appearance of gratings with and without the fundamental Fourier component. *J. Physiol. (Lond.)* **217**, 17–19.

Cohen, M. A. and Grossberg, S. (1984). Neural dynamics of brightness perception: Features, boundaries, diffusion and resonance. *Percept. Psychophys.* **36**, 428–456.

Cornsweet, T. (1970). *Visual Perception.* Academic Press, New York.

Dooley, R. P. and Greenfield, M. I. (1977). Measurements of edge-induced visual contrast and a spatial-frequency interaction of the Cornsweet illusion. *J. Opt. Soc. Am.* **67**, 761–765.

Du Buf, J. M. H. and Roufs, J. A. J. (1985). On brightness and brightness-contrast. Annual report of the Foundation of Biophysics, 228–236.

Foster, D. H. (1983). Experimental test of a network theory of vision. *Behavioural Brain Sci.* **6**, 664–665.

Georgeson, M. A. and Sullivan, G. D. (1975). Contrast constancy: deblurring in human vision by spatial frequency channels. *J. Physiol. (Lond.)* **252**, 627–656.

Grossberg, S. and Todorovic, D. (1988). Neural dynamics of 1-D and 2-D brightness perception: A unified model of classical and recent phenomena. *Percept. Psychophys.* **43**, 241–277.

Hamada, J. (1984). A multi-stage model for border contrast. *Biol. Cybernet.* **51**, 65–70.

Hamada, J. (1985). Asymmetric lightness cancellation in Craik–O'Brien patterns of negative and positive contrast. *Biol. Cybernet.* **52**, 117–122.

Hammond, P. (1974). Cat retinal ganglion cells: size and shape of receptive field centres. *J. Physiol.* **242**, 99–118.

Heinemann, E. G. (1972). Simultaneous brightness induction. In: *Handbook of Sensory Physiology VII/4.* D. Jameson and L. M. Hurvich (Eds). Springer-Verlag, New York.

Horn, B. K. P. (1974). Determining lightness from an image. *Comp. Graph. Image Process.* **3**, 277–299.

Hubel, D. H. and Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *J. Physiol. (Lond.)* **195**, 215–243.

Hurlbert, A. (1986). Formal connections between lightness algorithms. *J. Opt. Soc. Am. A* **3**, 1684–1693.

Isono, H. (1979). Measurement of edge-induced brightness. NHK Laboratories Note, Serial No. 233, 2–11.

Jory, M. and Day, R. H. (1979). The relationship between brightness contrast and illusory contours. *Perception* **8**, 3–9.

Kingdom, F. and Moulden, B. (1988). Border effects on brightness: A review of findings, models and issues. *Spatial Vision.* **3**, 225–262.

Kingdom, F. and Moulden, B. (1989). Modelling visual detection: luminance response nonlinearity and internal noise. *Q. J. Exp. Psychol.* **41A**, 675–696.

Kulikowski, J. J. (1976). Effective contrast constancy and linearity of contrast sensation. *Vision Res.* **16**, 1419–1431.

Laming, D. (1983). On the need for discipline in the construction of psychological theories. *Behav. Brain Sci.* **6**, 669–670.

Land, E. H. (1985). In: *Central and Peripheral Mechanisms in Colour Vision.* T. Ottosan and S. Zeki (Eds.). Macmillan, New York.

Land, E. H. and McCann, J. (1971) Lightness and retinex theory. *J. Opt. Soc. Am.* **61**, 1–11.

Legge, G. E. and Kersten, D. (1983). Light and dark bars: contrast discrimination. *Vision Res.* **23**, 473–483.

Magnusson, S. and Glad, A. (1975). Brightness and darkness enhancement during flicker: perceptual correlates of neuronal B- and D-systems in human vision. *Exp. Brain Res.* **22**, 399–413.

Morrone, M. C. and Burr, D. C. (1988). Feature detection in human vision: a phase dependent energy model. *Proc. R. Soc. Lond.* **B235**, 221–245.

Moulden, B. and Kingdom, F. (1989). White's effect: a dual mechanism. *Vision Res.* **29**, 1245–1259.

O'Brien, V. (1958). Contour perception, illusion and reality. *J. Opt. Soc. Am.* **48**, 112–119.

Reid, R. C. and Shapley, R. (1988). Brightness induction by local contrast and the spatial dependence of assimilation. *Vision Res.* **28**, 115–132.

Robson, J. G. (1975). Receptive fields: Neural representation of the spatial and intensive attributes of the visual image. In: *Handbook of Perception V: Seeing*, 81–117.

Schiller, P. H. (1986). The central visual system. *Vision Res.* **26**, 1351–1386.

Shapley, R. and Enroth-Cugell, C. (1984). Visual adaptation and retinal gain controls. In: *Progress in Retinal Research*. N. N. Osbourne and G. J. Chader (Eds.). Oxford, Pergamon Press, pp. 263–346.

Shevell, S. K. (1989). On neural signals that mediate induced blackness. *Vision Res.* **29**, 891–900.

Spillman, L. and Levine, J. (1971). Contrast enhancement in a Hermann grid with variable figure-ground ratio. *Exp. Brain Res.* **13**, 547–559.

Todorovic, D. (1987). The Craik–O'Brien–Cornweet effect: New varieties and their theoretical implications. *Percept. Psychophys.* **42**, 545–560.

Tolhurst, D. J. (1972). On the possible existence of edge detector neurons in human visual system. *Vision Res.* **12**, 797–804.

Van Erning, L. T. Th. D., Gerrits, J. M. and Eijkman, E. G. J. (1988). Apparent size and receptive field properties. *Vision Res.* **28**, 407–418.

Watt, R. J. (1988). *Visual Processes: Computational, Psychophysical and Cognitive Research*. Lawrence Erlbaum, Hove and London (UK).

Watt, R. J. and Morgan, M. J. (1985). A theory of the primitive spatial code in human vision. *Vision Res.* **11**, 1661–1674.

Whittle, P. (1986). Increments and decrements: luminance discrimination. *Vision Res.* **26**, 1677–1691.

Wilson, H. R. (1978). Quantitative characterisation of two types of line-spread function near the fovea. *Vision Res.* **18**, 971–981.

Wilson, H. R. (1980). A transducer function for threshold and suprathreshold human vision. *Biol. Cybernet.* **38**, 171–178.

Wilson, H. R. and Bergen, J. R. (1979). A four mechanism model for threshold vision. *Vision Res.* **19**, 19–32.